ARL-TR-7190 ● SEP 2015

US Army Research Laboratory

Characterization of Extremely Lightweight Intrusion Detection (ELIDe) Power Utilization by Varying N-gram and Hash Length

by Garrett S Payer, Ken F Yu, and Richard E Harang

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.

ARL-TR-7190 ● SEP 2015

US Army Research Laboratory

Characterization of Extremely Lightweight Intrusion Detection (ELIDe) Power Utilization by Varying N-gram and Hash Length

by Garrett S Payer, Ken F Yu, and Richard E Harang
Computational and Information Sciences Directorate, ARL

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED (From - To)
Sep 2015	Final	April 2014–September 2014

4. TITLE AND SUBTITLE
Characterization of Extremely Lightweight Intrusion Detection (ELIDe) Power Utilization by Varying N-gram and Hash Length

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Garrett S Payer, Ken F Yu, and Richard E Harang

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
US Army Research Laboratory
ATTN: RDRL-CIN-D
2800 Powder Mill Road
Adelphi, MD 20783-1138

8. PERFORMING ORGANIZATION REPORT NUMBER
ARL-TR-7190

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This document presents the results of the power-utilization and packet-loss study for the Extremely Lightweight Intrusion Detection (ELIDe) algorithm on an Android-based mobile device. Our results show that N-gram length has consistent power utilization using the range of 7–12 bits. However, if the hash length is higher than 13, it will have a drastic effect on the power utilization due to the resulting increase in computational time.

15. SUBJECT TERMS
ELIDe, Android

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Ken F Yu
Unclassified	Unclassified	Unclassified	UU	18	19b. TELEPHONE NUMBER (Include area code) 301-394-3181

Contents

List of Figures

1. Introduction

Signature detection is a very performance-intensive task in regard to network traffic.[1] Extremely lightweight intrusion detection (ELIDe) was developed to reduce the performance requirements for malicious-packet detection for use in low size, weight, and power (SWaP) devices such as mobile phones.[2] These devices can perform signature detection on network traffic in tactical environments where SWaP bandwidth, power, and battery life are limited.

For ELIDe to be useful on low-performance, low-power devices, its power utilization needs to be fully understood. In a related report,[3] we have evaluated the effect of ELIDe's power usage at varying throughput speeds to a mobile device. In addition, we evaluated how to vary the packet size while keeping throughput constant. The intention was to define a spectrum of possible traffic types to determine how they affect ELIDe's power utilization and performance.

Controlling the characteristics of network traffic as it is being sent to a mobile device is a very difficult task. In general, network traffic cannot be manipulated before being received by the device, unless a Berkley Packet Filter is used or the amount of network traffic that gets classified is limited.[4] However, the way in which ELIDe classifies packets can vary.

ELIDe reduces a network packet into a set of N-grams using a sliding window.[2] The length of these N-grams can be adjusted to be smaller or larger. In Chang et al.[2] the ELIDe was found to be most accurate when a 10-byte N-gram was used in evaluating the test data. However, classification of other types of malicious packets could be more or less accurate by increasing or decreasing the length of the N-gram used by ELIDe. The final weight used by ELIDe is still determined by the hash length.

As a packet is divided into N-grams, these features are then hashed and the hashes masked to a user-defined length.[2] The final hash length determines length of the final feature vector, as well as the length of the weight file that is created during the learning phase. The larger the hash length, the larger the feature vector—which means more calculations, and ultimately computational resources, required to compute the dot product between the weight and feature vector.

Both the N-gram and hash length are set during the learning phase.[2] Depending on the data set that the ELIDe is attempting to learn, these values can be varied in order to achieve greater accuracy. In addition, because hash length can greatly affect the performance required for the classification, the hash length can be set smaller to decrease the amount of processing and limit the amount of power utilized.

Depending on mission requirements, accuracy of detection can be lowered in favor of better performance and power utilization by adjusting the hash length.

To understand how varying these values affects power utilization, we have run a number of experiments to determine how N-gram and hash length impacts the power utilization on a mobile device.

2. Setup

Similar to the report on power utilization,[3] we utilized many of the same devices and the same test environment in order to characterize the power utilization of ELIDe. The ELIDe port to Android is a newer version that is modified to allow changing the weight file used without redeploying the application.

2.1 Mobile Device

We used the same Sprint-brand Galaxy S3 smart phone. The Galaxy S3 line of smart phones varied in its technical specifications depending on the carrier. For reference, the Sprint-brand Galaxy S3 has the following technical specifications:

- Qualcomm Snapdraggon S4 Plus MSM8960
- Dual core 1.5-GHz Krait Processor
- Adreno 225 graphics processor
- 2048 MB of random-access memory (RAM)
- 32-GB internal storage
- 2,100-mAh battery

2.2 Network

We utilize the mobile phone's Wi-Fi adapter for network traffic associated with a wireless access point. To reduce interference, the laptop generating network traffic utilizes an Ethernet port located on the access point. By limiting network connectivity through the Ethernet port, wireless traffic would be localized to the Wi-Fi communication between the mobile phone and the access point. (See Fig. 1.)

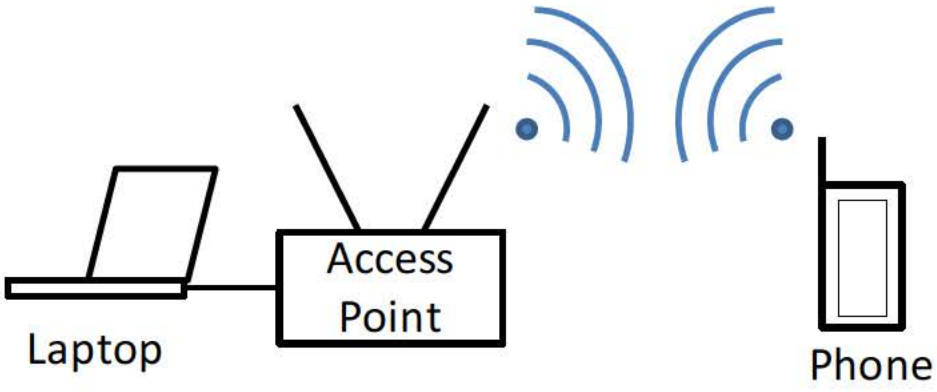

Fig. 1 Test-network setup

In addition, the phone was placed in "airplane mode" while leaving only the wireless network communication enabled. This prevented the device from 1) attempting to use any other wireless capabilities and 2) influencing the power readings. All notifications were also disabled during testing to prevent audio and vibration from interfering with the power-usage information.

The following devices were used alongside the mobile phone for this network:

- Cisco Wireless WAP300N
- Dell Latitude E6500–Core 2 Duo 2.53 GHz, 4 GB of RAM
- Kali Linux 1.0.7

2.3 Software Configuration

Similar to the work performed in our related report,[3] we utilized a version of the ELIDe software ported to Android for use on Android mobile devices. The software records the following statistics: packets classified, the number of packets that were positively and negatively classified, and the number of packets dropped after the buffer fills. The software also records battery statistics. Every time the available battery power dropped a percentage, the current runtime of the test was also recorded.

2.4 Determining Power Utilization

Similarly, too, was the need to determine the amount of additional power required to run ELIDe continuously.[3] By recording the ongoing runtime and battery percentage utilized, we can determine the amount of extra power used by ELIDe compared to when the device does not perform ELIDe classification. If the final runtime is significantly different between when ELIDe classification is used and

when it is not, we can determine that ELIDe classification will use significantly more power versus when the runtimes only vary a little.

To calculate the exact percentage of additional battery power that was used, the following formula is used:

$$\%\text{power} = \frac{\text{power consumption with ELIDe enabled}}{\text{power consumption without ELIDe}} \times 100$$

2.5 Experimental Setup

Using the test-network setup, our laptop was configured to send user datagram protocol[5] packets to the mobile device using the hping3 utility that is installed by default within the Kali Linux distribution. Because we are not varying our throughput or packet size, we sent packets with 600 bytes of padding at 1 megabit per second. The speed and padding used in the experiment remained constant throughout each test.

As noted in our related report,[3] the ELIDe Android software is configured to capture packets into a buffer, which will then be used to feed the classifier. In the event that packets are not processed fast enough and the buffer fills, packets will be dropped and escape classification.

We independently tested variations in both N-gram length and hash length. Because we were not testing throughput or packet size, the same control could be compared against each trial to determine the additional power required to run the classifier at different hash and N-gram lengths. The control data used the same throughput and packet padding; however, the classification was disabled so that the runtime would not be affected by increased power utilization of ELIDe classification.

3. Results

In the first part of our experiment, we measured the amount of additional battery power that was consumed when we changed the N-gram length during the initial testing phase. While a larger N-gram technically results in fewer N-grams being produced from a single network packet, in practice the number of N-grams that will be hashed will change very little between 2 lengths of several bytes' difference. Due to this, we did not expect the power utilization to significantly decrease as the N-gram length increased. However, we found that using different N-gram lengths to derive features from a packet did not consistently affect the amount of power utilized by ELIDe classification. (See Fig. 2.)

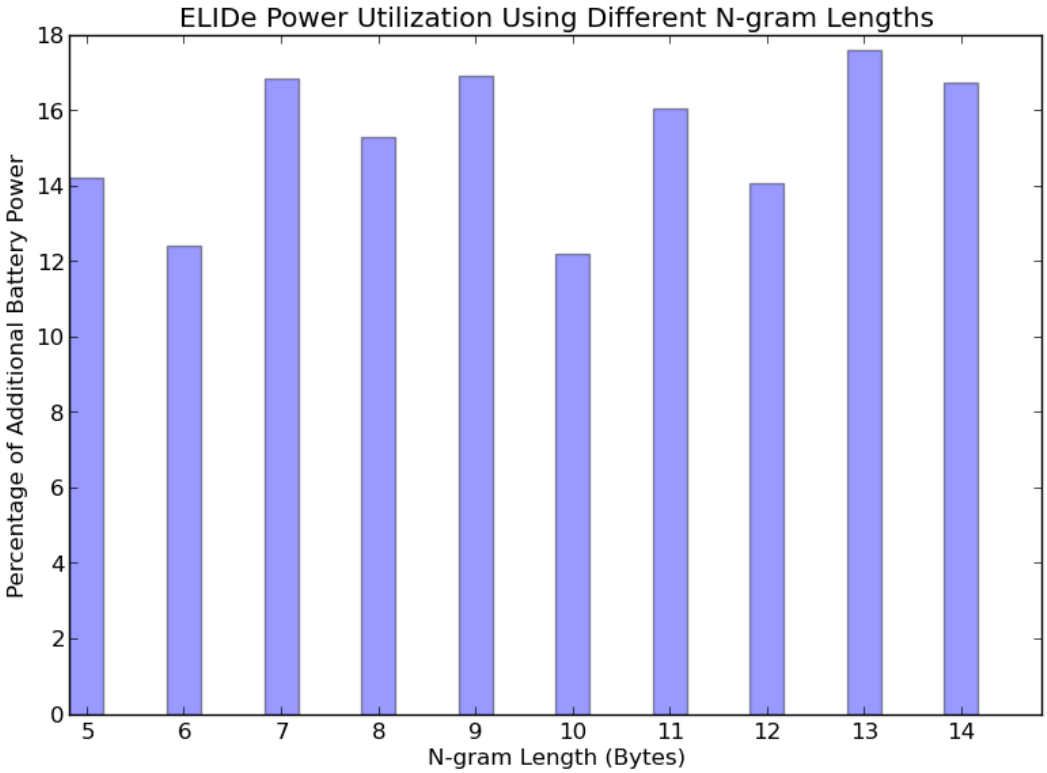

Fig. 2 Power utilization of ELIDe when varying N-gram length

We can analyze the graph and conclude that there is no direct relationship between power utilization and scaling the N-gram lengths. However, even-numbered lengths appear to generally use less power than their adjacent odd-numbered lengths, though the reason for this is unclear. A length of 10 bytes utilized the least amount of power in our experiments; however, this result may be due to random sampling error.

In previous experiments, we found that power utilization will decrease despite heavier loads due to dropped packets. This drop in power utilization would occur because the ELIDe would be unable to finish processing packets fast enough before the buffer filled, thereby dropping packets.[3] However, we found that packet loss was nearly nonexistent, as shown in Fig. 3.

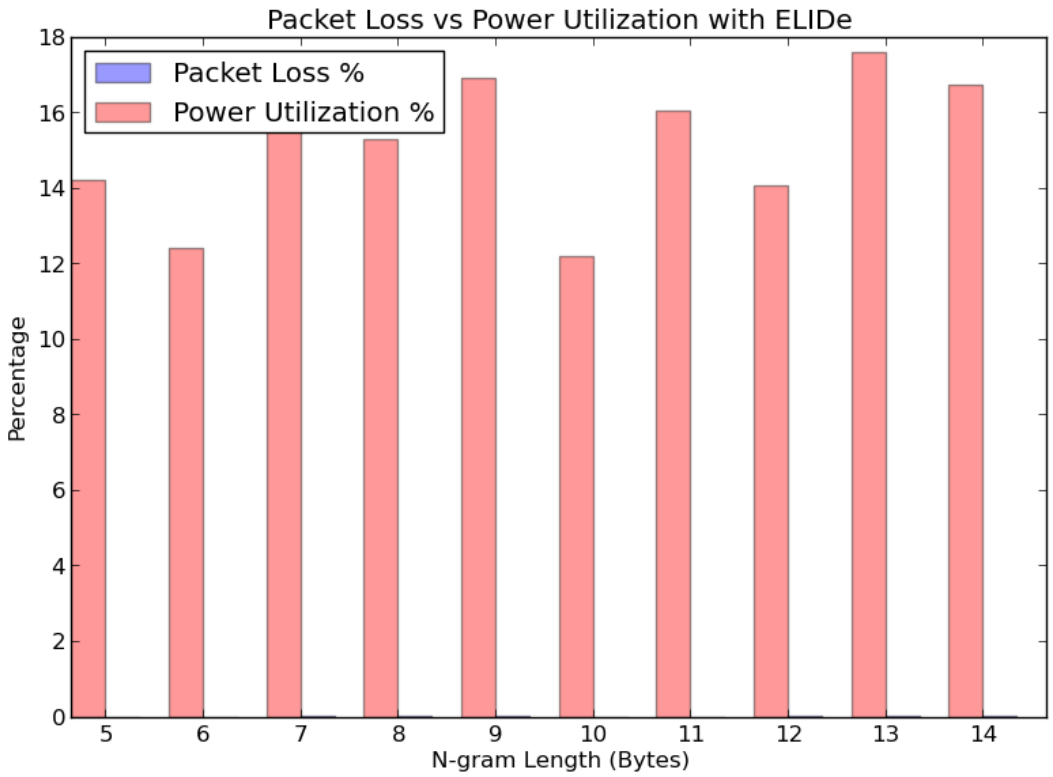

Fig. 3 Packet loss versus power usage of ELIDe when varying N-gram length

Insignificant amounts of packets were dropped during ELIDe classification; therefore, we conclude that varying the N-gram length does not affect packet loss, and thus would not skew power-usage data.

3.1 Hash-Bit Length

Although N-gram length can make a difference in classification accuracy, depending on the size of the features within a network packet, the length does not heavily influence the amount of power used by ELIDe classification. However, there is a relationship between the amount of calculations and the performance required when varying the length of the hash because hash length influences the feature-vector size. Note that if the hash-length width is increased by 1 bit, then the feature-vector size will double. Feature-vector size significantly impacts the number of calculations that need to be performed, which can lead to additional power utilization. We discovered that power utilization does not vary significantly at smaller hash sizes. (See Fig. 4.)

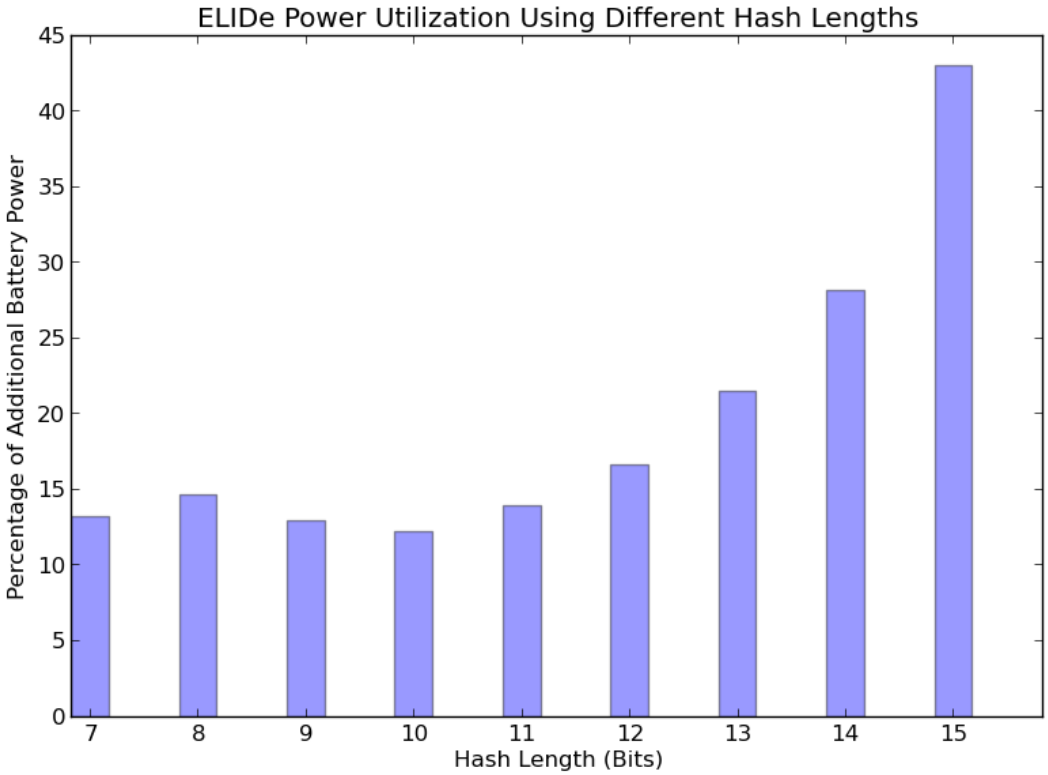

Fig. 4 Power utilization of ELIDe when using different hash lengths

Although it is not significantly different at lower hash lengths, power utilization begins to consistently increase at a length of 11 bits. The amount of performance required to utilize ELIDe is significantly higher than the performance needed for classification until at a hash length of 11 bits. Power utilization for hash lengths below 11 bits will not vary significantly; however, lengths larger than 11 bits require significantly more performance and power to perform ELIDe classification. We did not experience packet loss for the load that was put on the mobile device for all hash lengths except at a length of 15 bits (as seen in Fig. 5).

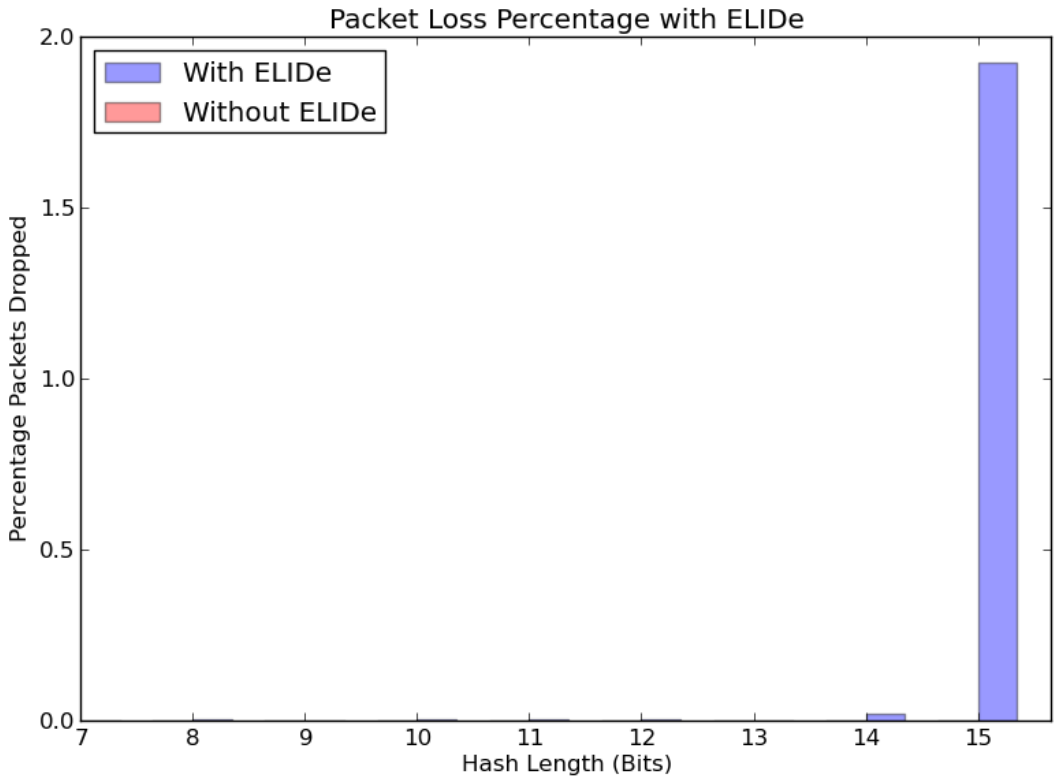

Fig. 5 Packet loss incurred at different hash lengths

The insignificant amount of power loss indicates that the throughput and packet size tested did not begin to exceed the performance capabilities of the mobile device using a 14-bit hash length, and packet loss only began to become significant at 15 bits.

We conclude that sacrificing the accuracy of longer hash lengths for shorter lengths will not decrease power utilization significantly when the hash length is less than 11 bits. However, lengths more than 11 bits significantly influence the amount of performance and thus battery utilization. The sacrifice of accuracy at these lengths will result in reduced power consumption.

4. Conclusions

The purpose of the work was to determine the effect of changing the N-gram and hash lengths to verify any difference in power utilization. A change in the N-gram length produced an inconsistent pattern in power utilization by ELIDe. Depending on the types of features that ELIDe can be trained to identify, changing the N-gram length may be prudent but will not significantly affect power utilization.

Changing the hash length used by ELIDe did produce a meaningful pattern. The larger the length, the larger the feature vector that ELIDe must utilize. The

increased computational load did lead to increased performance and battery utilization when longer than 11 bits, but it was not significantly different for lengths less than 11. As the length approached 15 bits, packets were dropped, indicating that the performance requirements were becoming higher than what the mobile device needed to keep up with the incoming packets. Thus, if the hash length needs to be adjusted to allow for better detection rates, power utilization would only be significantly impacted if the hash length starts to approach 15 bits in length.

5. References

1. Hugelshofer F, Smith P, Hutchison D, Race NJ. OpenLIDS: A lightweight intrusion detection system for wireless mesh networks. In: Proceedings of the 15th Annual International Conference on Mobile Computing and Networking, 2009 Sep 20–25; Beijing (China).

2. Chang RJ, Harang RE, Payer GS. Extremely Lightweight Intrusion Detection (ELIDe). Adelphi (MD): Army Research Laboratory (US); 2013 Dec. Report No.: ARL-CR-0730. Also available at http://www.dtic.mil/dtic/tr/fulltext/u2/a592893.pdf.

3. Payer GS, Yu K, Harang R. Characterization of ELIDe power utilization with varying throughput and payload sizes. Adelphi (MD): US Army Research Laboratory (US); 2015 Sep. Report No.: ARL-TN-0705.

4. McCanne S, Jacobson V. The BSD packet filter: a new architecture for user-level packet capture. In: Proceedings of the 1993 Winter USENIX Conference, 1993 Jan 25–29; San Diego (CA).

5. Postel J. User datagram protocol. Information Sciences Institute; RFC 768. 1980 Aug.

1 DEFENSE TECHNICAL
(PDF) INFORMATION CTR
DTIC OCA

2 DIRECTOR
(PDF) US ARMY RSRCH LAB
RDRL CIO LL
RDRL IMAL HRA RECORDS MGMT

1 DIRECTOR
(PDF) US ARMY RSRCH LAB
RDRLCIN D
K YU

INTENTIONALLY LEFT BLANK.